



US010846455B2

(12) **United States Patent**  
**Ghosh et al.**

(10) **Patent No.:** **US 10,846,455 B2**  
(45) **Date of Patent:** **Nov. 24, 2020**

(54) **AUTOMATIC DEFINITION AND EXTRACTION OF FUNCTIONAL COVERAGE METRIC FOR EMULATION-BASED VERIFICATION**

9/3869; G06F 13/00; G06F 13/22; G06F 13/4291; G06F 30/398; G06F 11/261; G06F 15/7864; G06F 16/211; G06F 16/245; G06F 1/06; G06F 2111/04; G06F 2111/20; G06F 2201/80; G06F 30/20; G06F 30/31; G06F 30/3308; G06F 30/337; G06F 30/34; G06F 30/367; G06F 30/373; G06F 30/396;

(71) Applicant: **Synopsys, Inc.**, Mountain View, CA (US)

(Continued)

(72) Inventors: **Saptarshi Ghosh**, Bangalore (IN); **Yogesh Pandey**, Bangalore (IN); **Sivaprasad Acharya**, Singapore (SG); **Eduard Cerny**, Marlborough, MA (US)

(56) **References Cited**

U.S. PATENT DOCUMENTS

(73) Assignee: **SYNOPSYS, INC.**, Mountain View, CA (US)

8,413,088 B1 \* 4/2013 Armbruster ..... G06F 30/33 716/106  
2005/0010880 A1 \* 1/2005 Schubert ..... G06F 30/33 716/106  
2008/0082946 A1 \* 4/2008 Zilic ..... G06F 30/3323 716/103

(\* ) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

\* cited by examiner

(21) Appl. No.: **16/297,324**

Primary Examiner — Binh C Tat

(22) Filed: **Mar. 8, 2019**

(74) Attorney, Agent, or Firm — Alston & Bird LLP

(65) **Prior Publication Data**

US 2019/0340327 A1 Nov. 7, 2019

(57) **ABSTRACT**

(30) **Foreign Application Priority Data**

Mar. 10, 2018 (IN) ..... 201811008880


A method of verifying a circuit design, includes, in part, identifying a first groups of signals associated with the circuit, selecting a signal sampling window depth, performing a first verification of the circuit using a first test bench adapted to cause transitions in the first group of signals, storing values of the signals in the first group during each of the cycles defined by the sapling window depth to generate a first functional coverage, performing a second verification of the circuit design using a second test bench to generate a second functional coverage, comparing the second functional coverage to the first functional coverage, and automatically generating one or more cover property statements if the second functional coverage is less than the first functional coverage. The one or more cover property statements cause the second functional coverage to become equal to or greater than the first functional coverage.

(51) **Int. Cl.**  
**G06F 30/20** (2020.01)  
**G06F 30/398** (2020.01)

(52) **U.S. Cl.**  
CPC ..... **G06F 30/398** (2020.01); **G06F 30/20** (2020.01)

(58) **Field of Classification Search**  
CPC ..... G06F 30/33; G06F 30/331; G06F 30/30; G06F 2115/10; G06F 2111/02; G06F 8/20; G06F 8/41; G06F 8/53; G06F 2117/08; G06F 2119/12; G06F 30/3312; G06F 30/3323; G06F 9/3867; G06F

**20 Claims, 3 Drawing Sheets**



(58) **Field of Classification Search**

CPC .... G01R 31/318357; G01R 31/318364; G01R  
31/318335; G01R 31/2848

USPC ..... 716/100-109

See application file for complete search history.




FIG. 1A




FIG. 1B




FIG. 1C




FIG. 2




FIG. 3

1

**AUTOMATIC DEFINITION AND  
EXTRACTION OF FUNCTIONAL  
COVERAGE METRIC FOR  
EMULATION-BASED VERIFICATION**

CROSS-REFERENCE TO RELATED  
APPLICATION

The present application claims benefit to Indian Patent Application Number 201811008880, filed Mar. 10, 2018, the content of which is incorporated herein by reference in its entirety.

The present application is related to commonly owned U.S. Pat. No. 9,626,468, the content of which is incorporated herein by reference in its entirety.

COPYRIGHT NOTICE

The assignee of this patent document has no objection to the facsimile reproduction by anyone of the patent document itself or of the patent application, as it appears in the files of the United States Patent and Trademark Office, but otherwise reserves all rights whatsoever in any included works of authorship protected by copyright.

DISCLAIMER

In the following Background, Summary, and Detailed Description, headings should not be construed as necessarily limiting. In the following Background, Summary and Detailed Description, the citation or identification of any publication does not signify relevance or status as prior art for any of the claimed or described embodiments. Paragraphs for which the text is all italicized signifies text that is common to multiple Synopsys patent specifications.

BACKGROUND

Design of modern semiconductors or other electronic devices is complex and often requires, in part, verification processes that are time consuming and costly. Often extensive analysis and verification methods are required to ensure that a circuit operates as intended before the circuit is manufactured. In order to perform such verifications, various simulation and other verification tools are used.

Traditional verification methods use test benches to apply a sequences of known test vectors to the circuit inputs and monitor the results at the circuit outputs. If the expected output values are received, the circuit is considered to function properly. However, the test bench approach, along with many currently implemented verification techniques, has become less efficient at finding design errors, both because of increased test duration and other difficulties inherent in verifying complex circuit designs. Further, current methods do not scale to large electronic systems that may include billions of transistors. A need continues to exist for improved circuit design verification techniques that provide maximize coverage while maintaining a low cost.

SUMMARY

The claims signify a brief description of one or more of the innovations, embodiments, and/or examples found within this disclosure.

A method of verifying a circuit design, in accordance with one embodiment of the present invention, includes, in part, identifying a first groups of signals associated with the

2

circuit, selecting a signal sampling window depth, performing a first verification of the circuit using a first test bench adapted to cause transitions in the first group of signals, storing, in response to the first test bench, values of the signals in the first group during each of the cycles defined by the sampling window depth to generate a first functional coverage, performing a second verification of the circuit design using a second test bench to generate a second functional coverage, comparing the second functional coverage to the first functional coverage, and generating, automatically, one or more cover property statements if the second functional coverage is less than the first functional coverage. The one or more cover property statements cause the second functional coverage to become equal to or greater than the first functional coverage.

In one embodiment, the first verification is a hardware emulation of the circuit design. In one embodiment, the second verification is a hardware emulation or a simulation of the circuit design. In one embodiment, the method further includes, in part, comprising storing values of the signals associated with the first functional coverage in a first memory having a storage space defined by a number of the signals in the group and the sampling window depth. In one embodiment, the method further includes, in part, setting a status bit in each address of the first memory containing the signal values.

In one embodiment, the method further includes, in part, storing values of the signals associated with the second functional coverage in a second memory having a storage space defined by a number of the signals in the group and the sampling window depth. In one embodiment, the method further includes, in part, setting a status bit in each address of the second memory containing the signal values. In one embodiment, the comparison is made by comparing the status bits of the first and second memories. In one embodiment, the first and second memories are defined in the same memory. In one embodiment, the memory is a register.

A computer system for verifying a circuit design, in accordance with one embodiment of the present invention, is configured to identify a first groups of signals associated with the circuit, select a signal sampling window depth, perform a first verification of the circuit using a first test bench adapted to cause transitions in the first group of signals, store, in response to the first test bench, values of the signals in the first group during each of the cycles defined by the sampling window depth to generate a first functional coverage, perform a second verification of the circuit design using a second test bench to generate a second functional coverage, compare the second functional coverage to the first functional coverage, and generate one or more cover property statements if the second functional coverage is less than the first functional coverage. The one or more cover property statements cause the second functional coverage to become equal to or greater than the first functional coverage.

In one embodiment, the first verification is a hardware emulation of the circuit design controlled by the computer. In one embodiment, the second verification is a hardware emulation or a simulation of the circuit design controlled by the computer. In one embodiment, the computer system is further configured to store values of the signals associated with the first functional coverage in a first memory having a storage space defined by a number of the signals in the group and the sampling window depth.

In one embodiment, the computer system is further configured to set a status bit in each address of the first memory containing the signal values. In one embodiment, the computer system is further configured to store values of the

signals associated with the second functional coverage in a second memory having a storage space defined by a number of the signals in the group and the sampling window depth.

In one embodiment, the computer system is further configured to set a status bit in each address of the second memory containing the signal values. In one embodiment, the computer system is further configured to perform the comparison by comparing the status bits of the first and second memories. In one embodiment, the first and second memories are defined in the same memory. In one embodiment, the memory is a register.

This Summary does not attempt to completely signify any particular innovation, embodiment, or example as it can be used in commerce. Additionally, this Summary is not intended to signify essential elements of an innovation, embodiment or example or to limit the scope of the subject matter of this disclosure.

The innovations, embodiments, and/or examples found within this disclosure are not all-inclusive, but rather describe the basic significance of the subject matter. Accordingly, one use of this Summary is as a prelude to a Detailed Description presented later.

### DRAWINGS

The following Detailed Description, Figures, appended Additional Figures and appended Claims signify the nature and advantages of the innovations, embodiments and/or examples of the claimed inventions. All of the Figures signify innovations, embodiments, and/or examples of the claimed inventions for purposes of illustration only and do not limit the scope of the claimed inventions. Such Figures are not necessarily drawn to scale, and are part of the Disclosure.

In the Figures, similar components or features may have the same, or similar, reference signs in the form of labels (such as alphanumeric symbols, e.g., reference numerals), and may signify similar or equivalent functionality. Further, various components of the same type may be distinguished by following the reference label by a dash and a second label that distinguishes among the similar components. If only the first reference label is used in the specification, the description is applicable to any one of the similar components having the same first reference label irrespective of the second reference label. A brief description of the Figures is below.

FIGS. 1A, 1B and 1C are simplified block diagrams of a computer system suitable for use with embodiments of the technology, as well as circuit design and circuit embodiments of the technology.

FIG. 2 illustrates various processes performed in the design, verification and fabrication of an integrated circuit design.

FIG. 3 is a block diagram of an emulation system.

In such various figures, reference signs may be omitted as is consistent with accepted engineering practice; however, one of ordinary skill in the art will understand that the illustrated components are readily understood when viewed in context of the illustration as a whole and the accompanying disclosure describing such various figures.

### DETAILED DESCRIPTION

The Figures and the following Detailed Description signify innovations, embodiments and/or examples by way of illustration only, with various features, structures or characteristics described together in a single embodiment to

streamline the disclosure. Variations of any of the elements, processes, machines, systems, manufactures or compositions disclosed by such exemplary innovations, embodiments and/or examples will be readily recognized and may be used in commerce without departing from the principles of what is claimed. The Figures and Detailed Description may also signify, implicitly or explicitly, advantages and improvements of a subset of the exemplary embodiments described herein.

In the Figures and Detailed Description, numerous specific details may be described to enable one or more of the exemplary innovations, embodiments and/or examples. In the interest of not obscuring the presentation of the exemplary innovations, embodiments and/or examples in the following Detailed Description, some processing steps or operations that are known in the art may be combined together for presentation and for illustration purposes and might not be described in detail. However, a person skilled in the art will recognize that these exemplary innovations, embodiments and/or examples may be used in commerce without these specific details or with equivalents thereof. In other instances, well-known processes and devices are not described in detail as not to unnecessarily obscure aspects of these exemplary innovations, embodiments and/or examples. In other instances, some processing steps or operations that are known in the art may not be described at all. Instead, the following description is focused on the distinctive features or elements of various exemplary innovations, embodiments and/or examples. Furthermore, while this description may refer to some components of the structure in the singular tense, more than one component may be depicted throughout the Figures and like components are labeled with like numerals.

In accordance with some embodiments of the present invention, the functional coverage of a verification run of a circuit design is increased via a set of statements to match the functional coverage of an earlier verification run. The earlier verification run may be achieved using hardware emulation, whereas subsequent verification runs may be achieved using computer simulation or hardware emulation of the circuit.

A circuit design may include millions of internal signals. Therefore, it is impractical to search in the signal space for arbitrary temporal relations over values of non-trivial subset of the signals. Likewise, in hardware emulation, any additional hardware instrumentation for coverage is costly. Therefore, line coverage and similar metrics are usually not deployed. That leaves assertion coverage or functional coverage as possible options.

Assertion depends on how well the user covers the design with assertion statements. Functional coverage, however, is more focused with transaction coverage than detailed functional coverage of the circuit design. Embodiments of the present invention mitigate the problem of coverage metrics that are either too costly or fail to address lower level functional coverage of a design by automatically extracting functional coverage points. To achieve this, embodiments of the present invention, in part, analyze the structure and behavior of the design to automatically generate assert property and cover property statements, as described further below.

Increasing the functional coverage points, in accordance with embodiments of the present invention is achieved in two phases, namely compilation and emulation/simulation. During compilation, a static analysis of the circuit design under verification (hereinafter alternatively referred to as design) is performed to identify groups of signals that are of

5

interest based on their connectivity. Such a signal group may include either primary inputs, outputs or internal registers that are in the same clock domain. Next, a sampling window of depth D, which may include one or more clock cycles, is selected. The signals together with the sampling window depth define the template for collecting the signal values during the emulation/simulation. For a group having N signals and a window depth D, there are (NxD) data points collected in every clock cycle, where x represents the multiplication symbol.

During instrumentation of the design, the signal values of interest in each signal group are stored in a memory allocated to that signal group. For a group having N signals and a window depth D representing the past values of such signals, a memory (which may be a register) with the storage space of  $2^{N \times D}$  bits is defined. The address to the memory is (NxD) bits wide.

During emulation (or simulation), the memory is initialized to 0. At every clock cycle the address is formed from the observed (NxD) values and a corresponding status bit in the memory is set to 1. At the end of the emulation (or simulation), the contents of all the signal group memories are collected for analysis.

Such instrumentation is performed for each signal group by instantiating the following exemplary module (described in Verilog) and connecting it to the N signals of the group:

```

module SAMPLE_SIG_GROUP
  #(ID=0, // id #, assigned sequentially
    N=1, // # of signals in group
    SEQ_DEPTH=1 // sampling depth, 1 means current values
  )
  (
    input CLK, // default sampling by posedge
    input RST_n, // reset on low
    input [N-1:0] IN
  );
  if (SEQ_DEPTH != inside{1:3})
    $error("Bad SEQ_DEPTH %0d", SEQ_DEPTH);
  reg [N-1:0] IN_CYCLE_1;
  reg [N-1:0] IN_CYCLE_2;
  always @(posedge CLK or negedge RST_n)
  if (!RST_n) begin
    IN_CYCLE_1 <= 0 ;
    IN_CYCLE_2 <= 0 ;
  end
end else begin

```

-continued

```

  if (SEQ_DEPTH==3) begin
    IN_CYCLE_1 <= IN;
    IN_CYCLE_2 <= IN_CYCLE_1;
  end
  else if (SEQ_DEPTH==2) begin
    IN_CYCLE_1 <= IN;
    IN_CYCLE_2 <= 0 ;
  end
end else begin
  IN_CYCLE_1 <= 0 ;

```

6

-continued

```

    IN_CYCLE_2 <= 0 ;
  end
end
reg MEM [0:(2**(N*SEQ_DEPTH)-1)]; // coverage memory
always @(posedge CLK or negedge RST_n)
if (!RST_n) begin
  // Don't add anything in memory while in reset
end
else begin
  // set memory bit at address obtained from sample
  if (SEQ_DEPTH==3) begin
    MEM[{IN_CYCLE_2,IN_CYCLE_1,IN}] <= 1'b1;
  end
  else if (SEQ_DEPTH==2) begin
    MEM[{IN_CYCLE_1,IN}] <= 1'b1;
  end
  else begin
    MEM[IN] <= 1'b1;
  end
end
end
endmodule

```

The design so instrumented is then compiled and synthesized for emulation or simulation. During emulation (or simulation), the memory is initialized to 0. At every clock cycle the address is formed from the observed (NxD) values and the corresponding status bit in the memory is set to 1. At the end of emulation (or simulation), the contents of all the signal group memories are collected for analysis. If the memory bit associated with any combination of the signals in a group is not set to 1, a cover property statement is automatically generated for that signal combination to increase the coverage.

Assume during a first hardware verification run (run A), a signal group containing variables y, en and d, respectively representing the output, the enable signal and the input signal of a latch, are defined. Assume a window depth of 2 is used for each of these signals. Therefore, the current values of each of these signals at clock cycle T, and their previous signal values at clock cycle T-1 is used. Accordingly, a sum of 5 values are sampled and stored during each clock cycle. Table 1 below shows exemplary values of each of these signals during 16 consecutive clock cycles:

TABLE 1

|    | Clk |   |   |   |   |   |   |   |   |    |    |    |    |    |    |    |
|----|-----|---|---|---|---|---|---|---|---|----|----|----|----|----|----|----|
|    | 1   | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 |
| en | 0   | 1 | 1 | 1 | 0 | 0 | 1 | 0 | 0 | 1  | 1  | 0  | 0  | 0  | 1  | 0  |
| d  | 1   | 0 | 0 | 1 | 1 | 0 | 0 | 1 | 0 | 1  | 0  | 0  | 0  | 1  | 1  | 1  |
| y  | 0   | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0  | 1  | 0  | 0  | 0  | 1  | 0  |

55

With the sampling window depth of 2 clock cycles, the observed signal values are obtained from Table 1, as shown in Table 2. Subscript 0 (e.g., do) refers to the current clock cycle value of a variable, and subscript -1 (e.g., d<sub>-1</sub>) refers to the immediately preceding clock cycle of the variable.

TABLE 2

| Row | en. | d <sub>-1</sub> | en <sub>0</sub> | d <sub>0</sub> | y <sub>0</sub> |
|-----|-----|-----------------|-----------------|----------------|----------------|
| 1   | 0   | 0               | 0               | 0              | 0              |

65

TABLE 2-continued

| Row | en. | d <sub>-1</sub> | en <sub>0</sub> | d <sub>0</sub> | y <sub>0</sub> |
|-----|-----|-----------------|-----------------|----------------|----------------|
| 2   | 0   | 0               | 0               | 1              | 0              |
| 3   | 0   | 0               | 1               | 0              | 0              |
| 4   | 0   | 0               | 1               | 1              | 0              |
| 5   | 0   | 1               | 0               | 0              | 0              |
| 6   | 0   | 1               | 1               | 0              | 1              |
| 7   | 0   | 1               | 1               | 1              | 1              |
| 8   | 1   | 0               | 0               | 0              | 0              |
| 9   | 1   | 0               | 0               | 1              | 0              |
| 10  | 1   | 0               | 1               | 0              | 0              |
| 11  | 1   | 0               | 1               | 1              | 0              |
| 12  | 1   | 1               | 0               | 1              | 0              |
| 13  | 1   | 1               | 1               | 0              | 1              |

For example, row number 2 of Table 2 is obtained from and corresponds to the values listed in clock cycles 13 and 14 of Table 1. To collect such data during, for example, an emulation phase, a memory size of 2<sup>5</sup>=32 bits is needed. At the beginning of the verification phase, all the memory spaces for the signal groups are initialized to 0. After generating the signal values shown above, the memory associated with the exemplary signal group will have stored a status bit of 1 (not shown in the Figures) at each of the 13 addresses shown in Table 2. Therefore, 13 of the 32 memory locations are covered for the above signal values.

Assume it is desired to achieve the same coverage during a second verification stage (run B), as was achieved during emulation run A. Assume that Table 3 shows the observed values of the same signal group during run B of the emulation or simulation.

TABLE 3

| Row | en. | d <sub>-1</sub> | en <sub>0</sub> | d <sub>0</sub> | y <sub>0</sub> |
|-----|-----|-----------------|-----------------|----------------|----------------|
| 1   | 0   | 0               | 0               | 0              | 0              |
| 2   | 0   | 0               | 0               | 1              | 0              |
| 3   | 0   | 0               | 1               | 0              | 0              |
| 4   | 0   | 0               | 1               | 1              | 0              |
| 5   | 0   | 1               | 0               | 0              | 0              |
| 6   | 0   | 1               | 1               | 0              | 1              |
| 7   | 0   | 1               | 1               | 1              | 1              |
| 8   | 1   | 0               | 0               | 0              | 0              |
| 9   | 1   | 0               | 0               | 1              | 0              |
| 10  | 1   | 0               | 1               | 0              | 0              |

It is seen that Table 3 is similar to Table 2 except that the signal values in rows 11-13 of run A are not observed in run B and thus are not present in Table 3. In other words, the coverage in run B is smaller than that in run A. To remedy this problem and ensure that run B has the same coverage as or greater coverage than run A, in accordance with one aspect of the present invention, the following three cover properties statements are generated (it is assumed that the clock is the same as that used for sampling the table):

- C1: cover property(en && !d ##1 en && d && !y);
- C2: cover property(en && d ##1 !en && d && !y);
- C3: cover property(en && d ##1 en && !d && y);

Cover property statement C1 generates the Boolean AND of signal en and inverse of signal d (i.e. !d) at clock cycle T-1, and the Boolean AND of signal en, signal d, and inverse of signal y (i.e., !y). Accordingly, cover property statement C1 generate for run B the signal values shown in rows number 11 of Table 2. In a similar manner, cover property statement C2 and C3 generate the signal values for run B that are shown in rows number 12 and 13 of Table 2, respectively. The cover property statement may be supplied to a coverage tool to generate the input sequences so as to

improve the test bench. Using the cover statements as described above, the functional coverage of run B becomes equal to that of run A.

The above-described embodiments may be used to develop a simulation testbench for a subsystem that is used in several full systems in emulation. The emulation runs can provide higher coverage which then can be reflected in a higher quality of the simulation run on the subsystem, with easier debug and analysis.

DETAILED DESCRIPTION—TECHNOLOGY SUPPORT

General Computer Explanation

FIGS. 1A, 1B and 1C are simplified block diagrams of a computer system suitable for use with embodiments of the technology, as well as circuit design and circuit embodiments of the technology.

In FIG. 1A, computer system 3110 typically includes at least one computer or processor 3114 which communicates with a number of peripheral devices via bus subsystem 3112. Typically, the computer can include, or the processor can be, any of a microprocessor, graphics processing unit, or digital signal processor, and their electronic processing equivalents, such as an Application Specific Integrated Circuit (ASIC) or Field Programmable Gate Array (FPGA). The terms ‘processor’ and ‘computer’ are further defined below. These peripheral devices may include a storage subsystem 3124, comprising a memory subsystem 3126 and a file storage subsystem 3128, user interface input devices 3122, user interface output devices 3120, and a network interface subsystem 3116. The input and output devices allow user interaction with computer system 3110.

The computer system may be a server computer, a client computer, a workstation, a mainframe, a personal computer (PC), a tablet PC, a set-top box (STB), a personal digital assistant (PDA), a cellular telephone, a smartphone, a web appliance, a rack-mounted “blade”, a kiosk, a television, a game station, a network router, switch or bridge, or any data processing machine capable of executing instructions 124 (sequential or otherwise) that specify actions to be taken by that machine.

The computer system typically includes an operating system, such as Microsoft’s Windows, Sun Microsystems’s Solaris, Apple Computer’s MacOS, Linux or Unix. The computer system also typically can include a Basic Input/Output System (BIOS) and processor firmware. The operating system, BIOS and firmware are used by the processor to control subsystems and interfaces connected to the processor. Typical processors compatible with these operating systems include the Pentium and Itanium from Intel, the Opteron and Athlon from Advanced Micro Devices, and the ARM processor from ARM Holdings.

Innovations, embodiments and/or examples of the claimed inventions are neither limited to conventional computer applications nor the programmable apparatus that run them. For example, the innovations, embodiments and/or examples of what is claimed can include an optical computer, quantum computer, analog computer, or the like. The computer system may be a multi-processor or multi-core system and may use or be implemented in a distributed or remote system. The term ‘processor’ here is used in the broadest sense to include a singular processor and multi-core or multi-processor arrays, including graphic processing units, digital signal processors, digital processors and combinations of these devices. Further, while only a single

computer system or a single machine may be illustrated, the use of a singular form of such terms shall also signify any collection of computer systems or machines that individually or jointly execute instructions 124 to perform any one or more of the sets of instructions discussed herein. Due to the ever-changing nature of computers and networks, the description of computer system 3110 depicted in FIG. 1A is intended only as one example for purposes of illustrating the preferred embodiments. Many other configurations of computer system 3110 are possible having more or less components than the computer system depicted in FIG. 1A.

Network interface subsystem 3116 provides an interface to outside networks, including an interface to communication network 3118, and is coupled via communication network 3118 to corresponding interface devices in other computer systems or machines. Communication network 3118 may comprise many interconnected computer systems, machines and communication links. These communication links may be wireline links, optical links, wireless links, or any other devices for communication of information. Communication network 3118 can be any suitable computer network, for example a wide area network such as the Internet, and/or a local area network such as Ethernet. The communication network can be wired and/or wireless, and the communication network can use encryption and decryption methods, such as is available with a virtual private network. The communication network uses one or more communications interfaces, which can receive data from, and transmit data to, other systems. Embodiments of communications interfaces typically include an Ethernet card, a modem (e.g., telephone, satellite, cable, or ISDN), (asynchronous) digital subscriber line (DSL) unit, Firewire interface, USB interface, and the like. One or more communications protocols can be used, such as HTTP, TCP/IP, RTP/RTSP, IPX and/or UDP.

User interface input devices 3122 may include an alphanumeric keyboard, a keypad, pointing devices such as a mouse, trackball, touchpad, stylus, or graphics tablet, a scanner, a touchscreen incorporated into the display, audio input devices such as voice recognition systems or microphones, eye-gaze recognition, brainwave pattern recognition, and other types of input devices. Such devices can be connected by wire or wirelessly to a computer system. In general, use of the term “input device” is intended to include all possible types of devices and ways to input information into computer system 3110 or onto communication network 3118. User interface input devices typically allow a user to select objects, icons, text and the like that appear on some types of user interface output devices, for example, a display subsystem.

User interface output devices 3120 may include a display subsystem, a printer, a fax machine, or non-visual displays such as audio output devices. The display subsystem may include a cathode ray tube (CRT), a flat-panel device such as a liquid crystal display (LCD), a projection device, or some other device for creating a visible image such as a virtual reality system. The display subsystem may also provide non-visual display such as via audio output or tactile output (e.g., vibrations) devices. In general, use of the term “output device” is intended to include all possible types of devices and ways to output information from computer system 3110 to the user or to another machine or computer system.

Memory subsystem 3126 typically includes a number of memories including a main random-access memory (RAM) 3130 (or other volatile storage device) for storage of instructions and data during program execution and a read only memory (ROM) 3132 in which fixed instructions are stored.

File storage subsystem 3128 provides persistent storage for program and data files, and may include a hard disk drive, a floppy disk drive along with associated removable media, a CD-ROM drive, an optical drive, a flash memory, or removable media cartridges. The databases and modules implementing the functionality of certain embodiments may be stored by file storage subsystem 3128.

Bus subsystem 3112 provides a device for letting the various components and subsystems of computer system 3110 communicate with each other as intended. Although bus subsystem 3112 is shown schematically as a single bus, alternative embodiments of the bus subsystem may use multiple busses. For example, RAM-based main memory can communicate directly with file storage systems using Direct Memory Access (DMA) systems.

FIG. 1B depicts a memory 3140 such as a non-transitory, computer readable data and information storage medium associated with file storage subsystem 3128, and/or with network interface subsystem 3116, and can include a data structure specifying a circuit design. The memory 3140 can be a hard disk, a floppy disk, a CD-ROM, an optical medium, removable media cartridge, or other medium that stores computer readable data in a volatile or non-volatile form. Software read into a computer from such a memory can be converted at a selected instance in time from a tangible form to a transmission signal that is propagated through a medium (such as a network, connector, wire, or trace as an electrical pulse or a medium such as space or an atmosphere as electromagnetic radiation with wavelengths in the electromagnetic spectrum longer than infrared light).

FIG. 1C signifies an integrated circuit 3190 created with the described technology that includes one or more cells selected, for example, from a cell library.

## DETAILED DESCRIPTION—TECHNOLOGY SUPPORT

### Hardware/Software Equivalence

Some of the innovations, embodiments and/or examples described herein comprise and/or use a processor. As used herein, the term ‘processor’ signifies a tangible data and information processing device that physically transforms data and information, typically using a sequence transformations (also referred to as ‘operations’). Data and information can be physically represented by an electrical, magnetic, optical or acoustical signal that is capable of being stored, accessed, transferred, combined, compared, or otherwise manipulated by the processor. The term “processor” can signify a singular processor and multi-core systems or multi-processor arrays, including graphic processing units, digital signal processors, digital processors or combinations of these elements.

The processor can be electronic, for example, comprising digital logic circuitry (for example, binary logic), or analog (for example, an operational amplifier). The processor can also be non-electronic, for example, as seen in processors based on optical signal processing, DNA transformations or quantum mechanics, or a combination of technologies, such as an optoelectronic processor. For data and information structured in binary form, any processor that can transform the data and information using the AND, OR and NOT logical operations (and their derivatives, such as the NAND, NOR, and XOR operations) can transform the data and information using any function of Boolean logic. A processor such as an analog neural network processor can also transform data and information non-digitally. There is no

scientific evidence that any of these processors are processing, storing and retrieving data and information, in any manner or form equivalent to the bioelectric structure of the human brain.

The one or more processors may also operate to support performance of the relevant operations in a ‘cloud computing’ environment or as a ‘software as a service’ (SaaS). For example, at least some of the operations may be performed by a group of processors available at a distributed or remote system, these processors accessible via a communications network (e.g., the Internet) and via one or more software interfaces (e.g., an application program interface (API).)

As used herein, the term ‘module’ signifies a tangible data and information processing device, that typically is limited in size and/or complexity. For example, the term ‘module’ can signify one or more methods or procedures that can transform data and information. The term ‘module’ can also signify a combination of one or more methods and procedures in a computer program. The term ‘module’ can also signify a small network of digital logic devices, in which interconnections of the logic devices give structure to the network. Methods and procedures comprising a module, specified in a specialized language, such as System C, can be used to generate a specification for a network of digital logic devices that process data and information with exactly the same results as are obtained from the methods and procedures.

A module can be permanently configured (e.g., hardwired to form hardware), temporarily configured (e.g., programmed with software), or a combination of the two configurations (for example, a structured ASIC). Permanently configured modules can be manufactured, for example, using Application Specific Integrated Circuits (ASICs) such as Arithmetic Logic Units (ALUs), Programmable Logic Arrays (PLAs), or Read Only Memories (ROMs), all of which are typically configured during manufacturing. Temporarily configured modules can be manufactured, for example, using Field Programmable Gate Arrays (FPGAs—for example, sold by Xilinx or Intel’s Altera), Random Access Memories (RAMs) or microprocessors. A module is configured to process data and information, typically using a sequence transformations (also referred to as ‘operations’) applied to the data and information (or in the case of ROMs and RAMS, transforming data and information by using the input information as an address for memory that stores output data and information), to perform aspects of the present innovations, embodiments and/or examples of the invention.

Modules that are temporarily configured need not be configured at any one instance in time. For example, a processor comprising one or more modules can have the modules configured at different times. The processor can comprise a set of one or more modules at one instance of time, and to comprise a different set of one or modules at a different instance of time. The decision to manufacture or implement a module in a permanently configured form, a temporarily configured form, or a combination of the two forms, may be driven by cost, time considerations, engineering constraints and/or specific design goals. The “substance” of a module’s processing is independent of the form in which it is manufactured or implemented.

As used herein, the term ‘algorithm’ signifies a process comprising a sequence or set of operations or instructions that a module can use to transform data and information to achieve a result. A module can comprise one or more algorithms. As used herein, the term ‘thread’ refers to a sequence of instructions that can comprise a subset of the

instructions of an entire process or algorithm. A process or algorithm can be partitioned into multiple threads that can be executed in parallel.

As used herein, the term ‘computer’ includes at least one information processor that, for example, can perform certain operations such as (but not limited to) the AND, OR and NOT logical operations using electronic gates that can comprise transistors, with the addition of memory (for example, memory based on flip-flops using the NOT-AND or NOT-OR operation). Such a processor is said to be Turing-complete or computationally universal. A computer, whether or not it is a digital computer, typically comprises many modules.

As used herein, the term ‘software’ or ‘program’ signifies one or more algorithms and data structures that configure a processor for use in the innovations, embodiments and examples described in this specification. Such devices configurable by software include one or more computers, for example, standalone, client or server computers, or one or more hardware modules, or systems of one or more such computers or modules. As used herein, the term “software application” signifies a set of data and instructions that configure the processor to achieve a specific result, for example, to perform word processing operations, or to encrypt a set of data.

As used herein, the term ‘programming language’ signifies a grammar and syntax for specifying sets of instruction and data that comprise software. Programming languages include assembler instructions, instruction-set-architecture (ISA) instructions, machine instructions, machine dependent instructions, microcode, firmware instructions, state-setting data, or either source code or object code written in any combination of one or more higher level languages, such as conventional procedural programming languages, for example, the “C” programming language or similar programming languages (such as SystemC), or object oriented programming language such as Smalltalk, C++ or the like, and any future equivalent programming languages.

Software is entered into, equivalently, read into, one or memories of the computer or computer system from a data and information storage device. The computer typically has a device for reading storage media that is used to transport the software or has an interface device that receives the software over a network. This process is discussed in the General Computer Explanation section.

#### Detailed Description—Technology Support

##### EDA System/Workflow Explanation

FIG. 2 illustrates various processes performed in the design, verification and fabrication of an item of manufacture such as an integrated circuit using software tools with a computer, and possibly special hardware-assisted tools, to transform and verify design data and instructions that represent the integrated circuit. These processes start with the generation of a product idea 3010 with information supplied by a designer, information which is transformed during a process to create an item of manufacture (referred to herein as a design or device) that uses an EDA software tool 3012, which may also be signified herein as EDA software, as a design tool, or a verification tool. When the design is finalized, it can be taped-out 3034, which typically is when artwork (e.g., geometric patterns) for the integrated circuit is sent to a fabrication facility to manufacture the mask set, which is then used to manufacture the integrated circuit. After tape-out, a semiconductor die is fabricated 3036 and

packaging and assembly processes **3018** are performed, which result in the finished integrated circuit **3020** which may also be signified herein as a circuit, device, component, chip or SoC (System on Chip).

Items of manufacture, for example, a circuit or system are used in commerce at a variety of levels of abstraction ranging from low-level transistor layouts to high-level description languages. Most designers start at high-level of abstraction to design their circuits and systems, using a hardware description language (HDL) such as VHDL, Verilog, SystemVerilog, SystemC, MyHDL or OpenVera. The high-level HDL is easier for developers to comprehend, especially for a vast system, and may describe highly complex concepts that are difficult to grasp using a lower level of abstraction. The HDL description may be converted into other levels of abstraction as is helpful to the developers. For example, a high-level description may be converted to a logic-level register transfer level (RTL) description, a gate-level (GL) description, a layout-level description, or a mask-level description. Each lower abstraction level introduces more detail into the design description. The lower-levels of abstraction may be generated automatically by computer, derived from a design library, or created by another design automation technique. An example of a specification language at a lower level of abstraction is SPICE, much used detailed descriptions of analog-based circuits.

A design process that uses an EDA software tool **3012** includes processes **3050-3068**, which are described below. This design flow description is for illustration purposes only and is not meant to limit the present disclosure. For example, an actual integrated circuit design may require a designer to perform the design operations in a different sequence than the sequence described herein.

During system design **3050**, a designer describes the functionality to be manufactured. The designer can also perform what-if planning to refine the functionality and to check costs, etc. Hardware-software architecture partitioning can occur at this stage. Exemplary EDA software products from Synopsys, Inc. of Mountain View, Calif., that can be used at this stage include: Model Architect, Saber, System Studio, and Designware products.

During logic design and functional verification **3052**, modules in the circuit are specified in one or more hardware description languages, or HDLs, and the design in HDL is checked for functional accuracy, that is, to match the requirements of the specification of the circuit or system being designed to ensure that the design produces the correct outputs. Exemplary HDL languages are Verilog, VHDL and SystemC. Functional verification is typically done by using software-based simulators and other tools such as testbench generators, static HDL checking tools and formal verification tools. In some situations, special hardware referred to as 'emulators' or 'prototyping systems' are used to speed up the functional verification. Exemplary EDA software products from Synopsys, Inc. of Mountain View, Calif. that can be used at this stage include: VCS, Vera, Designware, Magellan, Formality, ESP and Leda products. Exemplary emulator and prototyping products also available from Synopsys that can be used at this state include: Zebu® and Protolink® (RTM="Registered Trademark").

During synthesis and design for test **3054**, HDL code is translated to a netlist. This netlist can be optimized for the target technology. Additionally, tests can be designed and implemented to check the finished integrated circuit. Exemplary EDA software products from Synopsys, Inc. of Mountain View, Calif. that can be used at this stage include:

Design Compiler, Physical Compiler, Test Compiler, Power Compiler, FPGA Compiler, Tetramax, and Designware products.

During netlist verification **3056**, the netlist is checked for compliance with timing constraints and for correspondence with the HDL code. Exemplary EDA software products from Synopsys, Inc. of Mountain View, Calif. that can be used at this stage include: Formality, Primetime, and VCS products.

During design planning **3058**, an overall floor plan for the integrated circuit is constructed and analyzed for timing and top-level routing. Exemplary EDA software products from Synopsys, Inc. of Mountain View, Calif. that can be used at this stage include: Astro and IC Compiler products.

During layout implementation **3060**, the physical placement (positioning of circuit elements such as transistors or capacitors) and routing (connection of the same by a plurality of conductors) occurs, as can selection of library cells to perform specific logic functions. Exemplary EDA software products from Synopsys, Inc. of Mountain View, Calif. that can be used at this stage include: the Astro and IC Compiler products.

During analysis and extraction **3062**, the circuit function is verified at the layout level, which permits refinement of the layout design. Exemplary EDA software products from Synopsys, Inc. of Mountain View, Calif. that can be used at this stage include: Astrorail, Primerail, Primetime, and Star RC/XT products.

During physical verification **3064**, the layout design is checked to ensure correctness for manufacturing constraints such as DRC constraints, electrical constraints, lithographic constraints, and circuitry function matching the HDL design specification. Exemplary EDA software products from Synopsys, Inc. of Mountain View, Calif. that can be used at this stage include the Hercules product.

During resolution enhancement **3066**, geometric manipulations of the layout are performed to improve manufacturability of the design. Exemplary EDA software products from Synopsys, Inc. of Mountain View, Calif. that can be used at this stage include: Proteus products.

During tape-out, data is created to be used (after lithographic enhancements are applied if appropriate) for production of lithography masks. Example EDA software products from Synopsys, Inc., that can be used during tape-out include the IC Compiler and Custom Designer families of products.

During mask-data preparation **3068**, the 'tape-out' data is used to produce lithography masks that are used to produce finished integrated circuits. Exemplary EDA software products from Synopsys, Inc. of Mountain View, Calif. that can be used at this stage include the CATS family of products.

For all of the above mentioned integrated circuit design tools, similar tools from other EDA vendors, such as Cadence, Siemens, other corporate entities or various non-commercial tools from universities, or open source repositories, can be used as an alternative.

Embodiments of the present disclosure can be used during one or more of the above-described stages. Specifically, some embodiments of the present disclosure can be used in EDA software **3012**.

A storage subsystem is preferably used to store the programs and data structures that provide the functionality of some or all of the EDA tools described herein, and tools applied for development of cells for the library and for physical and logical design using the library. These programs and data structures are generally executed by one or more processors in a manner known to those of ordinary skill in the art.

DETAILED DESCRIPTION—TECHNOLOGY  
SUPPORT

## Emulation Environment Explanation

An EDA software system, such as element **3012** depicted in FIG. 3, typically includes an emulation system **3054** to verify the functionality of the circuit design. FIG. 8 depicts a typical emulation system which includes a host computer system **3300** (often part of an EDA system) and an emulator system **3202** (typically a set of programmable devices such as Field Programmable Gate Arrays (FPGAs)). The host system generates data and information, typically using a compiler **3210**, to configure the emulator to emulate a circuit design. One of more circuit designs to be emulated are referred to as a DUT (Design Under Test). The emulator is a hardware system that emulates a DUT, for example, to use the emulation results for verifying the functionality of the DUT. One example of an emulation system that can be used for the embodiments disclosed herein is the ZeBu Server available from Synopsys, Inc.

The host system **3300** comprises one or more processors. In the embodiment where the host system is comprised of multiple processors, the functions described herein as being performed by the host system may be distributed among the multiple processors.

The host system **3300** typically includes a compiler **3210** that processes code written in a hardware description language that represents a DUT, producing data (typically binary) and information that is used to configure the emulation system **3202** to emulate the DUT. The compiler **3210** may transform, change, reconfigure, add new functions to, and/or control the timing of the DUT.

The host system and emulator exchange data and information using signals carried by an emulation connection. The connection can be one or more electrical cables, for example, cables with pin configurations compatible with the RS232 or USB protocols. The connection can be a wired communication medium or network, such as a local area network, or a wide area network such as the Internet. The connection can be a wireless communication medium or a network with one or more points of access, using a wireless protocol such as Bluetooth® or IEEE 802.11. The host system and emulator can exchange data and information through a third device, such as a network server.

The emulator includes multiple FPGAs (or other programmable devices), for example, elements **3204<sub>1</sub>** to **304<sub>N</sub>** in FIG. 8. Each FPGA can include one or more FPGA interfaces through which the FPGA is connected to other FPGAs of the emulator (and potentially other emulator hardware components), in order for the FPGAs to exchange signals. An FPGA interface may also be referred to as an input/output pin or an FPGA pad. While some embodiments disclosed herein make use of emulators comprising FPGAs, other embodiments can include other types of logic blocks instead of, or along with, the FPGAs for emulating DUTs, for example, custom FPGAs, specialized ASICs for emulation or prototyping, memories, and input/output devices.

A programmable device may include an array of programmable logic blocks and a hierarchy of reconfigurable interconnects that allow the programmable logic blocks to be connected to each other according to the descriptions in the HDL code. Each of the programmable logic blocks can be configured to perform complex combinational functions, or merely simple logic gates, such as AND, and XOR logic blocks.

In many FPGAs, logic blocks also include memory elements, which may be simple latches, flip-flops or more complex blocks of memory. Depending on the length of the interconnections between different logic blocks, signals may arrive at input terminals of the logic blocks at different times.

Programmable processors **3204<sub>1</sub>-3204<sub>N</sub>** may be placed into one or more hardware boards **3212<sub>1</sub>** through **3212<sub>m</sub>**. Many of such boards may be placed into a hardware unit, e.g. **3214<sub>1</sub>**. The boards within a unit may be connected using the backplane of the unit or any other types of connections. In addition, multiple hardware units (e.g., **3214<sub>1</sub>** through **3214<sub>k</sub>**) may be connected to each other by cables or any other means to form a multi-unit system. In general, the hardware emulation or prototype system **3202** may be formed using a single board, a single unit with multiple boards, or with multiple units without departing from the teachings of the present disclosure.

For a DUT that is to be emulated, the emulator receives from the host system one or more bit files including a description of the DUT. The bit files further specify partitions of the DUT created by the host system with trace and injection logic, mappings of the partitions to the FPGAs of the emulator, and design constraints. Based on the bit files, the emulator configures the FPGAs to perform the functions of the DUT. With some emulators, one or more FPGAs of an emulator already have the trace and injection logic built into the silicon of the FPGA. For this type of emulator, the FPGAs don't have to be configured by the host system to emulate trace and injection logic.

The host system **3110** receives (e.g., from a user) a description of a DUT that is to be emulated. In one embodiment, the DUT description is in a hardware description language (HDL), such as register transfer language (RTL). In another embodiment, the DUT description is in netlist level files, or a mix of netlist level files and HDL files. If part of the DUT description or the entire DUT description is in a HDL, the host system synthesizes the DUT description to create a gate level netlist based on the DUT description. A host system can use the netlist of the DUT to partition the DUT into multiple partitions, with some of these partitions including trace and injection logic. The trace and injection logic traces interface signals exchanged via the interfaces of an FPGA. Additionally, the trace and injection logic can be used to inject traced interface signals into the logic of the FPGA. The host system maps each partition to an FPGA of the emulator. With some emulators, the trace and injection logic is only included in select partitions for a group of FPGAs. The trace and injection logic can be built into one or more of the FPGAs of an emulator. The host system can synthesize multiplexers to be mapped into the FPGAs. The multiplexers can be used by the trace and injection logic to inject interface signals into the DUT logic.

The host system creates bit files describing each partition of the DUT and the mapping of the partitions to the FPGAs. For partitions in which trace and injection logic was incorporated, the bit files also describe the incorporation of the logic. The bit files may also include place and route information and design constraints. The host system stores the bit files and also stores for components of the DUT information describing which FPGAs are to emulate each component of the DUT (to which FPGAs each component is mapped).

Upon request, the host system transmits the bit files to the emulator. The host system instructs the emulator to emulate the DUT. During emulation of the DUT or at the end of the emulation, the host system receives emulation results from the emulator through the emulation connection. Emulation

results are data and information generated by the emulator based on the emulation of the DUT. The emulation results include interface signals (states of interface signals) traced by the trace and injection logic of each FPGA. The host system can store the emulation results, or transmit them to another processing system.

After emulation of the DUT, a user may request to debug a component of the DUT. If such a request is made the user may provide a time period of the emulation to debug. The host system identifies which FPGAs are configured to emulate the component based on the stored information. The host system retrieves stored interface signals associated with the time period and traced by the trace and injection logic of each identified FPGA. The host system instructs the emulator to re-emulate the identified FPGAs, either one by one, multiple at a time, or altogether. The host system transmits the retrieved interface signals to the emulator in order to re-emulate the component for the time period. The trace and injection logic of each identified FPGA injects its respective interface signals received from the host system into the logic of the DUT mapped to the FPGA. In case of multiple re-emulations of an FPGA, the results may be merged all together to have a full debug view.

The host system receives from the emulator signals traced by logic of the identified FPGAs during the re-emulation of the component. The host system stores the signals received from the emulator. The signals traced during the re-emulation can have a higher sampling rate than during the initial emulation. For example, in the initial run a traced signal may be comprised of a saved hardware state every X milliseconds. However, in the re-emulation the traced signal may be comprised of a saved hardware state every Y milliseconds, where Y is less than X. If the user requests to view a waveform of a signal traced during the re-emulation, the host system can retrieve the stored signal and display a plot of the signal (generates a waveform of the signal). Afterwards the user can, for example, request to re-emulate the same component but for a different time period or to re-emulate another component.

A host system typically comprises at least seven sub-systems: a design synthesizer, a mapping module, a run time module, a results module, a debug module, a waveform module, and a storage module. Each of these sub-systems may be embodied as hardware, software, firmware, or a combination thereof. Together these components configure the emulator, and monitor the emulation results.

The design synthesizer converts the HDL of a DUT into gate level logic. For a DUT that is to be emulated, the design synthesizer receives a description of the DUT. If the description of the DUT is fully or partially in HDL (e.g., RTL or other level of abstraction), the design synthesizer synthesizes the HDL of the DUT to create a gate-level netlist with a description of the DUT in terms of gate level logic.

The mapping module partitions DUTs and maps partitions to emulator FPGAs. The mapping module partitions a DUT at the gate level into a number of partitions using the netlist of the DUT. For each partition, the mapping module retrieves a gate level description of the trace and injection logic and incorporates the logic into the partition. As described above, the trace and injection logic included in a partition is configured to trace signals exchanged via the interfaces of an FPGA to which the partition is mapped (trace interface signals). The trace and injection logic can be incorporated into the DUT prior to the partitioning. For example, the trace and injection logic may be incorporated by the design synthesizer prior to or after the synthesizing the HDL of the DUT. Hence, the trace and injection logic

may not match the partitions, it may be a subset, a superset or even different from the partitions.

In addition to including the trace and injection logic, the mapping module may include additional tracing logic in a partition in order to trace the states of certain DUT components that are not traced by the trace and injection logic (to trace signals other than the interface signals traced by the trace and injection logic). The mapping module may include the additional tracing logic in the DUT prior to the partitioning or in partitions after the partitioning. The design synthesizer can include the additional tracing logic in an HDL description of the DUT prior to synthesizing the description.

The mapping module maps each partition of the DUT to an FPGA of the emulator. The mapping module performs the partitioning and mapping using design rules, design constraints (e.g., timing or logic constraints), and information about the emulator. For components of the DUT, the mapping module stores information in the storage module describing which FPGAs are to emulate each component.

Based on the partitioning and the mapping, the mapping module generates one or more bit files that describe the created partitions and the mapping of logic to each FPGA of the emulator. The bit files may include additional information, such as constraints of the DUT, and routing information of connections between FPGAs and connections within each FPGA. The mapping module can generate a bit file for each partition of the DUT, which can be stored in the storage module. Upon request from a user, the mapping module transmits the bit files to the emulator, which the emulator uses to configure the FPGAs to emulate the DUT.

If the emulator includes specialized ASICs that include the trace and injection logic, the mapping module may generate a specific configuration allowing to connect them to the DUT or just save the information of what traced/injected signal is and where the information is stored on the specialized ASIC.

The run time module controls emulations performed on the emulator. The run time module may cause the emulator to start or stop executing an emulation. Additionally, the run time module may provide input signals/data to the emulator. The input signals may be provided directly to the emulator through the connection or indirectly through other input signal devices. For example, the host system with the run time module may control an input signal device to provide the input signals to the emulator. The input signal device may be, for example, a test board (directly or through cables), signal generator, another emulator, or another host system.

The results module processes emulation results generated by the emulator. During emulation and/or after completing the emulation, the results module receives emulation results from the emulator generated during the emulation. The emulation results include signals traced during the emulation. Specifically, the emulation results include interface signals traced by the trace and injection logic emulated by each FPGA. The emulation results may also include signals traced by additional logic included in the DUT. Each traced signal can span multiple cycles of the emulation. A traced signal is comprised of multiple hardware states and each hardware state is associated with a time of the emulation. The results module stores the traced signals received in the storage module. For each stored signal, the results module can store information indicating which FPGA generated the traced signal.

The debug module allows users to debug DUT components. After the emulator has emulated a DUT and the results

module has received the interface signals traced by the trace and injection logic during the emulation, a user may request to debug a component of the DUT by re-emulating the component for a specific time period. In a request to debug a component, the user identifies the component and indicates a time period of the emulation to debug. The user's request can also include a sampling rate that indicates how often hardware states should be saved by logic that traces signals.

The debug module identifies the one or more FPGAs of the emulator that are configured to emulate the component based on the information stored by the mapping module in the storage module. For each identified FPGA, the debug module retrieves, from the storage module, interface signals traced by the trace and injection logic of the FPGA during the time period indicated by the user (i.e., retrieve hardware states traced by the trace and injection logic that are associated with the time period).

The debug module transmits the retrieved interface signals to the emulator. The debug module instructs the debug module to run the identified FPGAs and for the trace and injection logic of each identified FPGA to inject its respective traced signals into logic of the FPGA in order to re-emulate the component for the requested time period. The debug module can also transmit the sampling rate provided by the user to the emulator so that the tracing logic traces hardware states at the proper intervals.

To debug the component, the emulator only has to run the FPGAs to which the component has been mapped. Additionally, the re-emulation of the component does not have to start from the beginning, but can start at any point desired by the user.

For an identified FPGA, the debug module can transmit instructions to the emulator to load multiple emulator FPGAs with the same configuration of the identified FPGA. The debug module additionally instructs the emulator to run the multiple FPGAs in parallel. Each FPGA from the multiple FPGAs is run with a different time window of the interface signals in order to generate a larger time window in a shorter amount of time. For example, for the identified FPGA to run a certain amount of cycles it may take an hour. However, if multiple FPGAs are loaded with the configuration of the identified FPGA and each of the FPGAs runs a subset of the cycles, it may only take a few minutes for the FPGAs to collectively run all of the cycles.

A user may identify a hierarchy or a list of DUT signals to re-emulate. To enable this, the debug module determines the FPGA needed to emulate the hierarchy or list of signals, retrieves the necessary interface signals and transmits the retrieved interface signals to the emulator for re-emulation. Hence, a user can identify any element (e.g., component or signal) of the DUT to debug/re-emulate.

The waveform module generates waveforms based on traced signals. If a user requests to view a waveform of a signal traced during an emulation run, the host system retrieves the signal from the storage module. The waveform module displays a plot of the signal to the user. For one or more signals, when the signals are received from the emulator, the waveform module can automatically generate the plots of the signals.

#### DETAILED DESCRIPTION—SEMANTIC SUPPORT

The signifier 'commercial solution' signifies, solely for the following paragraph, an engineering domain-specific (and thus non-preemptive—see Bilski) electronic structure, process for specified machines, manufacturable circuit (and

their Church-Turing equivalents) or composition of matter that is useful in commerce to solve a problem of technology, that is, a use in commerce of an application of science or use in commerce of technology.

The signifier 'abstract' (when used in a patent claim for any embodiments disclosed herein for a new commercial solution that is a scientific application of one or more laws of nature {see Benson}, and that solves a problem of technology {see Diehr} used in commerce—or improves upon an existing commercial solution {see Diehr})—is precisely defined by the inventor(s) {see MPEP 2111.01} as follows:

a) a new commercial solution is 'abstract' if it is not novel (e.g., it is so well known in equal prior art {see Alice} and/or the use of equivalent prior art solutions is long prevalent {see Bilski} in science, engineering or commerce), and thus unpatentable under 35 U.S.C. 102, for example, because it is "difficult to understand" {see Merriam-Webster definition for 'abstract'} how the commercial solution differs from equivalent prior art solutions; or

b) a new commercial solution is 'abstract' if it is obvious, that is, if the existing prior art includes at least one analogous prior art solution {see KSR}, or the existing prior art includes at least two prior art items that can be combined {see Alice} by a person having ordinary skill in the art {a "PHOSITA", see MPEP 2141-2144} to be equivalent to the new commercial solution, and is thus unpatentable under 35 U.S.C. 103, for example, because it is "difficult to understand" how the new commercial solution differs from a PHOSITA-combination/-application of the existing prior art; or

c) a new commercial solution is 'abstract' if it is not disclosed with an enabling description either because there is insufficient guidance in the enabling description, or because only a generic implementation is described {see Mayo} with unspecified elements, parameters or functionality, so that a PHOSITA is unable to instantiate a useful embodiment of the new commercial solution, without, for example, requiring special programming {see Katz} or circuit design to be performed by the PHOSITA), and is thus unpatentable under 35 U.S.C. 112, for example, because it is "difficult to understand" how to use in commerce any embodiment of the new commercial solution.

#### DETAILED DESCRIPTION—CONCLUSION

The foregoing Detailed Description signifies in isolation the individual features, structures, functions, or characteristics described herein and any combination of two or more such features, structures, functions or characteristics, to the extent that such features, structures, functions or characteristics or combinations thereof are based on the present specification as a whole in light of the knowledge of a person skilled in the art, irrespective of whether such features, structures, functions or characteristics, or combinations thereof solve any problems disclosed herein, and without limitation to the scope of the claims. When an embodiment of a claimed invention comprises a particular feature, structure, function or characteristic, it is within the knowledge of a person skilled in the art to use such feature, structure, function, or characteristic in connection with other embodiments whether or not explicitly described, for example, as a substitute for another feature, structure, function or characteristic.

In view of the foregoing Detailed Description it will be evident to a person skilled in the art that many variations may be made within the scope of innovations, embodiments

and/or examples, such as function and arrangement of elements, described herein without departing from the principles described herein. One or more elements of an embodiment may be substituted for one or more elements in another embodiment, as will be apparent to those skilled in the art. The embodiments described herein are chosen to signify the principles of the invention and its useful application, thereby enabling others skilled in the art to understand how various embodiments and variations are suited to the particular uses signified.

The foregoing Detailed Description of innovations, embodiments, and/or examples of the claimed inventions has been provided for the purposes of illustration and description. It is not intended to be exhaustive nor to limit the claimed inventions to the precise forms described but is to be accorded the widest scope consistent with the principles and features disclosed herein. Obviously, many variations will be recognized by a person skilled in this art. Without limitation, any and all equivalents described, signified or incorporated by reference in this patent application are specifically incorporated by reference into the description herein of the innovations, embodiments and/or examples. In addition, any and all variations described, signified or incorporated by reference herein with respect to any one embodiment are also to be considered taught with respect to all other embodiments. Any such variations include both currently known variations as well as future variations, for example any element used herein includes a future equivalent element that provides the same function, regardless of the structure of the future equivalent.

It is intended that the scope of the claimed inventions be defined and judged by the following claims and equivalents. The following claims are hereby incorporated into the Detailed Description, with each claim standing on its own as a separate embodiment. Disclosed embodiments can be described with more features than are expressly recited in the claims.

What is claimed is:

1. A method of verifying a circuit design, the method comprising:  
 identifying a first group of signals associated with the circuit design;  
 selecting a signal sampling window depth;  
 performing a first verification of the circuit design using a first test bench adapted to cause transitions in the first group of signals;  
 storing, in response to the first test bench, values of the signals in the first group during each of the cycles defined by the signal sampling window depth to generate a first functional coverage;  
 performing a second verification of the circuit design using a second test bench to generate a second functional coverage;  
 comparing the second functional coverage to the first functional coverage;  
 determining a difference in the first functional coverage from the second functional coverage, wherein the difference in the first functional coverage from the second functional coverage comprises a subset of the values of the signals used in the generation of the first functional coverage; and  
 generating, automatically, one or more cover property statements, each of the one or more cover property statements being associated with a portion of the difference; and  
 providing the one or more cover property statements to a test bench to generate a new input sequence.

2. The method of claim 1 wherein the first verification is a hardware emulation of the circuit design.

3. The method of claim 2 wherein the second verification is a hardware emulation or a simulation of the circuit design.

4. The method of claim 1 further comprising storing values of the signals associated with the first functional coverage in a first memory having a storage space defined by a number of the signals in the group and the sampling window depth.

5. The method of claim 4 further comprising setting a status bit in each address of the first memory containing the signal values.

6. The method of claim 5 further comprising storing values of the signals associated with the second functional coverage in a second memory having a storage space defined by a number of the signals in the group and the sampling window depth.

7. The method of claim 6 further comprising setting a status bit in each address of the second memory containing the signal values.

8. The method of claim 7 wherein said comparison is made by comparing the status bits of the first and second memories.

9. The method of claim 8 wherein said first and second memories are defined in the same memory.

10. The method of claim 9 wherein said memory is a register.

11. A computer system for verifying a circuit design, the computer system configured to:

identify a first group of signals associated with the circuit design;  
 select a signal sampling window depth;  
 perform a first verification of the circuit design using a first test bench adapted to cause transitions in the first group of signals;  
 store, in response to the first test bench, values of the signals in the first group during each of the cycles defined by the sampling window depth to generate a first functional coverage;  
 perform a second verification of the circuit design using a second test bench to generate a second functional coverage;  
 compare the second functional coverage to the first functional coverage;  
 determine a difference in the first functional coverage from the second functional coverage, wherein the difference in the first functional coverage from the second functional coverage comprises a subset of the values of the signals used in the generation of the first functional coverage; and  
 generate one or more cover property statements, each of the one or more cover property statements being associated with a portion of the difference; and  
 providing the one or more cover property statements to a test bench to generate a new input sequence.

12. The computer system claim 11 wherein the first verification is a hardware emulation of the circuit design controlled by the computer.

13. The computer system of claim 12 wherein the second verification is a hardware emulation or a simulation of the circuit design controlled by the computer.

14. The computer system claim 11 wherein the computer system is further configured to:

store values of the signals associated with the first functional coverage in a first memory having a storage space defined by a number of the signals in the group and the sampling window depth.

15. The computer system of claim 14 wherein said computer system is further configured to set a status bit in each address of the first memory containing the signal values.

16. The computer system of claim 15 wherein said 5 computer system is further configured to store values of the signals associated with the second functional coverage in a second memory having a storage space defined by a number of the signals in the group and the sampling window depth.

17. The computer system of claim 16 wherein said 10 computer system is further configured to set a status bit in each address of the second memory containing the signal values.

18. The computer system of claim 17 wherein said computer system is further configured to perform the com- 15 parison by comparing the status bits of the first and second memories.

19. The computer system of claim 18 wherein said first and second memories are defined in the same memory.

20. The computer system of claim 19 wherein said 20 memory is a register.

\* \* \* \* \*